

**Best
Available
Copy**

AD-A008 865

FINAL TECHNICAL REPORT, OCTOBER 11, 1974
"THE ALOHA SYSTEM"

Norman Abramson

Hawaii University

Prepared for:

National Aeronautics and Space Administration
Advanced Research Projects Agency

11 October 1974

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE

ADA008865

128134

FINAL TECHNICAL REPORT
FOR CONTRACT NUMBER NAS2-6700

OCTOBER 11, 1974

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
U.S. Department of Commerce
Springfield, VA. 22151

THE ALOHA SYSTEM

UNIVERSITY OF HAWAII
HONOLULU, HAWAII 96822

DDC
RECEIVED
MAY 6 1975
D

(53)

FINAL TECHNICAL REPORT

October 11, 1974

ARPA Order Number: 1956

Program Code Number: 2P10

Name of Contractor: University of Hawaii

Effective Date of Contract: November 1, 1971

Contract Expiration Date: October 11, 1974


Amount of Contract: \$2,571,266

Contract Number: NAS2-6700

Principal Investigator and Phone Number: Norman Abramson (808) 948-7490

Project Scientist or Engineer and Phone Number: Norman Abramson (808) 948-7490

Short Title of Work: THE ALOHA SYSTEM



Sponsored by
Advanced Research Projects Agency
ARPA Order No. 1956

ia

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

SUMMARY

Research in THE ALOHA SYSTEM under Contract Number NAS2-6700 is divided into two major tasks; (1) to study and develop advanced forms of computer-communications networks using random-access packet switching methods (Task I), and (2) to conduct general studies of multiprocessor system organization centered on the development of the BCC 500 computer (Task II).

I. TASK I

1.0 OBJECTIVES

The work of Task I concerned the study and development of THE ALOHA SYSTEM and its extension to advanced forms of computer communications networks. Its objectives were to perform theoretical studies on and experimental tests of radio linked ALOHA type networks. The principal subtasks of this part were:

- (a) To perform theoretical and simulation studies of ALOHA type radio channels for use in packet switched communications networks.
- (b) To experimentally test improved versions of the ALOHA communications techniques and its extensions. A principal subtask was to design and develop a packet radio repeater suitable for use with THE ALOHA SYSTEM operational network.

Both of these subtasks were successfully completed and the results obtained are explained in detail in section I of this report. In addition, the results have been reported in detail in numerous external publications, ALOHA SYSTEM reports and internal documents. Task I publications for the period covered by Contract NAS2-6700 are given at the end of this section (4.0).

2.0 BACKGROUND

Developments in remote access computing during the latter part of the 1960's have resulted in increasing importance of remote time-sharing, remote job entry and networking for large information processing systems. The present generation of computer-communication systems is based on the use of leased or dial-up common carrier facilities, primarily wire connections. Under many conditions such communication facilities offer the best possible communications option to the overall system designer of a large computer-communication facility. In other circumstances, however, the organization of common carrier data communications systems seriously limits the design of a large information processing system.

When the constraint of data communications by wire is eliminated a number of options for different methods of organizing data communications within a computer-communications net are made available to the system designer. THE ALOHA SYSTEM project has investigated the use of new forms of random access communications for computer-communication networks; the first links in THE ALOHA SYSTEM UHF radio-linked computer system were set up in 1971.

Since that time THE ALOHA SYSTEM has been in continuous operation. The ALOHANET uses two channels at 407.350 MHz and at 413.475 MHz in the UHF band. These channels are now operated at 9600 baud. ALOHA uses packet communication techniques similar to that employed by the ARPANET in conjunction with a novel form of random-access radio-channel multiplexing. The radio channel, when used in a burst random-access mode, has come to be known as an ALOHA channel. The ALOHA technique, first described in a paper published in 1970, has opened up the entire field of packet broadcasting of which the ALOHANET, the ARPANET Satellite System, and the Packet Radio System are three examples.

Perhaps the most significant and promising application of the ALOHA technique is in satellite packet broadcasting. The design of the ARPA Satellite System uses a variation of the pure ALOHA technique to make feasible the sharing of a single satellite channel or transponder among a large number of users on a random access basis. The satellite channel can be regarded as a resource which can be shared by many users, and it is this resource-sharing that promises to significantly lower the cost of data transmission when the new U.S. domestic satellites become operational in the next few years.

The satellite broadcasting field is still new and largely untested. The satellite IMP (SIMP) developed by Bolt, Beranek and Newman is scheduled to go into field operation sometime in 1974. At present, however, THE ALOHA SYSTEM has the only burst-mode packet broadcasting satellite channel in operation (employing the NASA satellite ATS-1).

During the three years that THE ALOHA SYSTEM has been supported by ARPA through Contract NAS2-6700, Task I has developed an operational ground radio network; an important integrated version of the terminal control unit (TCU); an advanced communications controller/multiplexer including a network "gateway"; a packet repeater; a packet broadcasting satellite system using ATS-1 and a multiprocessor system simulation facility. During this period, ALOHA's contribution in the theoretical areas have been highlighted by a considerable number of ARPANET Satellite System Notes and studies for the ARPA Packet Radio System.

The motivation for this system building has been (1) to demonstrate the feasibility of the ALOHA multiplexing technique in a variety of situations, and (2) to gain insight at a practical level into the problems of the design and use of such a system. We believe that innovations in the architecture of computer-communication networks we have developed and the results obtained in our theoretical investigations have amply justified this work.

3.0 SPECIFIC ACCOMPLISHMENTS

3.1 Theoretical Results

Researchers in THE ALOHA SYSTEM have provided much of the impetus for the theoretical interest in multi-user channels which now exists in the U.S. The first multiple user channel designed specifically for computer-communication networks was the ALOHA channel and a good deal of work has gone into an examination of the properties of this channel under Contract NAS2-6700. Among the results obtained were those dealing with different data rate users and how they affect the channel in a multi-access mode. Results were obtained for both the unslotted case (where it was shown that a mixture of packet lengths tend to degrade the channel) and for the slotted case (where it was shown that a mixture of users with different transmission probabilities could lead to higher efficiency -- or "excess capacity"). Both of these results were important in obtaining an understanding of this new form of data communications.

Additional theoretical results were obtained dealing with the spatial capacity of an ALOHA channel, or how such a channel performs in the presence of strong signal capture effects. In Packet Radio Temporary Note 49, we analyzed the reception of packets in the presence of capture. We were able to show that when the central receiver is surrounded by a uniform density population of packet transmitting terminals, there exists a radius r_0 such that any packet transmitted from a terminal located within r_0 will be received correctly with probability one (after a finite number of retransmissions) while the expected number of retransmissions required for a packet transmitted from a terminal further from the center than r_0 will be unbounded. Thus there exists a natural circle of radius r_0 such that terminals transmitted from within this circle can get their packets into the central receiver,

while terminals transmitting from outside this circle spend all of their time retransmitting their packets in vain. We have called r_0 the *Sisyphus* distance of the ALOHA channel. N.T. Gaarder extended the Sisyphus result to the case of two sinks, both with perfect capture. He found that the capacity of that system was twice the capacity of a single terminal, and furthermore, it was independent of the actual throughput distribution. Thus the two ALOHA terminals automatically distribute the throughput so that each operates at peak efficiency.

Ronald Ibaraki, under the guidance of Professor Gaarder, produced a report "A Dynamic Analysis of ALOHA Systems with Blocking and Carrier Sense" analyzing some simple properties of dynamic ALOHA channels for the first time. From his mathematical model he obtained a system of first order differential equations which describe the dynamic behavior of an ALOHA channel. In this model attempts to send new packets and blocked packets form a Poisson point process. From the solutions, Ibaraki found that the throughput may approach $1/\tau$ but in the operation condition, the average delay approaches $N\tau$ and the users are blocked most of the time. By reducing throughput, the average delay and average number of blocked users can be reduced.

3.2 System Development and Experiments

The principal Task I subtask in this part of THE ALOHA SYSTEM contract during the period covered by this report was the design and development of an ALOHA repeater suitable for use with our two frequency random access communication network. Since the transmission scheme of the operational ALOHANET is by line-of-sight, repeaters are necessary to operate in an environment such as Honolulu where a diversity of terrain (mountains, high-rise buildings, heavy foliage, etc.) severely limits the radio range. An ALOHA repeater was initially built for testing purposes with simple hard-wired logic, similar to

the integrated ALOHA TCU's built in 1972.

Preliminary tests on Oahu indicated that the repeater functioned as designed and in January 1974, tests were made successfully from Mount Haleakala on Maui to test the range of the repeater and the effects of radio interference upon its operation. After the radio portion was tested on this repeater a more versatile programmable ALOHA repeater was constructed using an INTEL 8080 microprocessor chip. The repeaters now available will allow the investigation of routing algorithms with multiple repeater paths for a simple network of the sort in operation.

Another project which was strongly emphasized during the period of this contract was the design of an integrated or programmable terminal control unit (PCU) which is the major communication module at the terminal end of the ALOHANET. Existing hard-wire versions of the terminal control unit (TCU) were felt to be too bulky and expensive. Thus we concentrated on the development of two new TCU's involving an INTEL 8008 microcomputer and then an INTEL 8080 microcomputer on a single integrated circuit chip. The PCU enabled the system to respond to a variety of different transmission protocols, including variable length packets and character-by-character transmission. Two versions of the PCU were completed during the period covered by this report and the software has been successfully debugged. The job is an ongoing project and many of the experimental protocols possible with a PCU are yet to be implemented.

The establishment of the Lockheed SUE minicomputer facility was another system development of the ALOHANET during the period of this contract. This stand-alone computer was successfully integrated into the existing radio network and put into use as a stand-alone computing facility, an ALOHA simulation facility and a source of file traffic for the network.

In addition to work on the ALOHANET, satellite experiments in packet broadcasting have been conducted on the NASA ATS-1 satellite. The initial experiments were conducted with NASA designed modems and interface hardware. The equipment supplied by NASA/AMES (other than the radio) consists primarily of a PCM Bit Synchronizer, a NASA-built Formatter/Synchronizer, and a Convolutional Coder/Decoder. This equipment, taken together, represents an investment in the order of \$10,000,00. By going to a block code and implementing the error correction function in software, the cost of the stand alone coder/decoder could be saved. A standard 9.6K bps ALOHA modem was modified for rapid burst acquisition and low false-alarm rate. Minor modifications were made to the ATS-1 radio to interface the ALOHA modem and a simple interface unit was built to go between the modem and the radio. Tests to date have indicated the ALOHA modem arrangement to have about the same error rate capability as the NASA designed system. Therefore, the use of the cheaper ALOHA system appears quite feasible. The cost of the ALOHA modem should be no more than \$500.00

Efforts were directed toward investigation of the ATS-1 channel characteristics, development of bit error analysis software, of terminal acknowledgement capability in the MENEHUNE for use on ATS-1, and investigations of an improved modem design for use on the ATS-1 channel.

Measurements of bit-errors per packet and packet throughput using both the ALOHA and the NASA modems at various data rates were conducted. Preliminary measurements indicated that the ALOHA modem had superior packet throughput at a 9600 bps rate in comparison to the NASA-supplied equipment. No comparisons are available at other bit rates since the ALOHA modem was not operated at any other bit rate. In general, we found the ALOHA modem to have a throughput of 80 to 90 percent of all its packets without any errors.

These measurements were taken with the satellite operating in the full power mode. With the satellite in the low power mode (6 dB less) the throughput was seriously degraded for both systems to the order of about 10 percent packet throughput without errors. The above measurements were made on the basis of receiving a signal level of about one microvolt at the receiver preamp with the satellite in the full power mode.

In addition to the hardware system development reported above, considerable effort during the period of this contract was devoted to software development. Work was completed on the NCP and TELNET programs required to interface THE ALOHA SYSTEM to ARPANET.

The ALOHA-NCP program which supports ARPANET connections has been operational since June 1974. It is presently limited to half-duplex line-by-line terminal support, which has limited its use by project personnel. Character-by-character support is expected to be in operation at a later date for newer (programmable) TCU's. The use of ARPANET hosts with ALOHA terminals also created some problems which did not exist when using the UH360, such as the need for carriage return padding and a more sophisticated method of screen-size control for CRT displays. Solutions to these problems were subsequently implemented in the MENEHUNE. Additions to the NCP program to support Server-TELNET connections and provide appropriate flow control for file transfers were made and detailed documentation was completed for the present NCP version.

During this period a number of significant additions were made to the original ALOHA SYSTEM protocols. These consisted of new packet formats to support character-by-character and variable length packets, provisions for text transparency and flow-control signalling, and special protocols to accomplish automatic loading of programmable TCU's. The latter were required for support of SRI's portable TCU to be tested with our system, and also for

use with the INTEL 8080 TCU's. These protocols have been implemented in the MENEHUNE and documented. A significant amount of effort was also expended in MENEHUNE programming support of ATS-1.

4.0 DOCUMENTATION FOR PERIOD NOVEMBER 1, 1971 TO SEPTEMBER 30, 1974

BOOKS

- [1] ABRAMSON, NORMAN and KUO, FRANKLIN F. (co-editors): Computer-Communication Networks, Prentice-Hall Publishing Company, 1973.

PAPERS and ARTICLES

- [2] ABRAMSON, NORMAN: "The ALOHA System," Computer-Communication Networks, N. Abramson and F.F. Kuo, Editors, Prentice-Hall Publishing Company, 1973, pp. 501-518; also published as ALOHA System Technical Report B72-1, University of Hawaii, January 1972.
- [3] ABRAMSON, NORMAN and KUO, FRANKLIN F.: "Some Advances in Radio Communications for Computers," Seventh Annual IEEE Computer Society International Conference, COMPCON'73 Digest of Papers, February 1973, pp. 57-60; also published as Packet Radio Temporary Note 10, NIC Document 13643, Stanford Research Institute, California, January 1973; and as ALOHA System Technical Report B73-1, University of Hawaii, March 1973.
- [4] ABRAMSON, NORMAN: "Packet Switching with Satellites," AFIPS Conference Proceedings, National Computer Conference, Vol. 42, New York, June 1973, pp. 695-702; reprinted in Advances in Computer Communications, Wesley W. Chu, Editor, ARTECH House, Inc., 1974, pp. 68-75; also published as ALOHA System Technical Report B73-2, University of Hawaii, March 1973; and as ARPANET Satellite System Note 37, NIC Document 15194, Stanford Research Institute, California, March 1973.
- [5] LIN, SHU: "Multifold Euclidean Geometry Codes," IEEE Transaction on Information Theory, Vol. IT-19, No. 4, July 1973, pp. 537-548; also published as ALOHA System Technical Report A73-1, University of Hawaii, March 1973.
- [6] KUO, FRANKLIN F. and BINDER, RICHARD: "Computer-Communications by Radio and Satellite: The ALOHA System," Proceedings of the NATO Advanced Study Institute on Computer Communication Networks, F.F. Kuo and R.L. Grimsdale, co-editors, Noordhoff International Publishing Co., September 1973; also published as ALOHA System Technical Report B73-4, University of Hawaii, August 1973.
- [7] KUO, FRANKLIN F. and BINDER, RICHARD: "The Use of a Multi-Processor Minicomputer for Communication System Simulation," Proceedings of the Seventh Asilomar Conference on Circuits, Systems and Computers, Western Publishing Co., November 1973; pp. 313-319; also published as ALOHA System Technical Report B73-6, University of Hawaii, November 1973; and as Packet Radio Temporary Note 84, NIC Document 20397, Stanford Research Institute, California, November 1973.

- [8] ABRAMSON, NORMAN: "Digital Broadcasting in Hawaii -- The ALOHA System," *EDUCOM Bulletin*, Vol. 9, No. 1, Spring 1974, pp. 9-13; also published as *ALOHA System Technical Report B73-5*, University of Hawaii, November 1973.
- [9] FUJITA, K.; IKEDA, T.; NOGUCHI, S.; SATO, R.; OIZUMI, J.; EBIHARA, Y.; KUO, F.F.; and ABRAMSON, N.: "A Japan-Hawaii Computer Net -- Telex and Satellite," *Proceedings of the Seventh Hawaii International Conference on System Sciences - Subconference on Computer Nets*, Western Periodicals Co., January 1974, pp. 78-80.
- [10] KUO, FRANKLIN F.: "The ALOHA System," *Computer Communications Review*, Vol. 4, No. 1, January 1974, pp. 5-8.
- [11] KUO, FRANKLIN F.: "The ALOHA Broadcast Packet Communications System," *Proceedings of the NORTHWEST 74 CIPS-ACM Pacific Regional Symposium*, Vancouver, Canada, 23-24 May 1974, pp. 459-465; also published in the *Proceedings of the IRIA International Workshop on Modelling and Measurement of Computer Architectures and Networks*, France, 14-16 August 1974, pp. 275-283.
- [12] KUO, FRANKLIN F.: "On Packet Computer-Communication Networks," *Proceedings of the IEE 1974 European Conference on Circuit Theory and Design*, London, England, 23-26 July 1974, pp. 389-394.
- [13] KUO, FRANKLIN F.: "Political and Economic Issues for Internetwork Connection," *Proceedings of the 1974 International Conference on Computer Communications*, Stockholm, Sweden, 12-14 August 1974, pp. 389-391; also published as *International Network Working Group Note 58, NIC Document 30506*, Stanford Research Institute, California, 2 May 1974.

TECHNICAL REPORTS

- [14] CASE, RICHARD H.: "Gomoku," *ALOHA System Technical Report A71-9*, University of Hawaii, December 1971.
- [15] YAMAMOTO, RICHARD: "Computational Accuracy of Linear Stationary Iterations," *ALOHA System Technical Report A71-10*, University of Hawaii, December 1971.
- [16] BASS, CHARLIE C. and NICHOLSON, VINCENT J.: "Using Interactive Systems as a Tool in Creating Interactive Subsystems," *ALOHA System Technical Report A72-1*, University of Hawaii, February 1972.
- [17] LEE, WRENWICK: "The ALOHA System Sixteen-Channel Multiplexor Program Module Description," *ALOHA System Technical Report B72-2*, University of Hawaii, May 1972.
- [18] DAVIDSON, JOHN: "The ALOHA System Interface to TSO," *ALOHA System Technical Report B72-3*, University of Hawaii, May 1972.
- [19] LIAO, HENRY HERNG-JIUNN: "Multiple Access Channels," *ALOHA System Technical Report A72-2*, University of Hawaii, September 1972.

- [20] BASS, CHARLIE, C.: "Design and Implementation of UHTSS -- The University of Hawaii Time-Sharing System," *ALOHA System Technical Report B72-4*, University of Hawaii, September 1972.
- [21] STOUTENMYER, DAVID R.: "Simplifying Transformations for Multiply -- Symmetric Systems that are Asymmetrically Loaded," *ALOHA System Technical Report A73-2*, University of Hawaii, March 1973.
- [22] ABRAMSON, NORMAN and AH MAI, KAREN: "Pacific Educational Computer Network Study," *ALOHA System Technical Report B73-3*, University of Hawaii, May 1973.
- [23] PATEL, BHADRA: "Optimization of Large Scale Systems," *ALOHA System Technical Report A73-3*, University of Hawaii, December 1973.
- [24] KUO, FRANKLIN F. (Editor): "Annual Technical Report (December 31, 1973)," *ALOHA System Technical Report B74-1*, University of Hawaii, January 1974.
- [25] CHATTERGY, RAHUL: "Cost-Delay Invariance in Optimal Computer-Communication Networks," *ALOHA System Technical Report A74-1*, University of Hawaii, January 1974.
- [26] IBARAKI, RONALD Y.: "A Dynamic Analysis of ALOHA Systems with Blocking and Carrier Sense," *ALOHA System Technical Report B74-2*, University of Hawaii, January 1974; also published as *Packet Radio Temporary Note 87*, NIC Document 20407, Stanford Research Institute, California, January 1974.
- [27] STOUTENMYER, DAVID R.: "Automatic Error Analysis Using the Computer Symbolic Manipulation Language, Reduce," *ALOHA System Technical Report A74-2*, University of Hawaii, January 1974.
- [28] KUO, FRANKLIN F.: "Selected Papers on Computer Communication Nets," *ALOHA System Technical Report B74-4*, University of Hawaii, May 1974.
- [29] LU, SHYUE-CHING: "A Comparison of Multi-Access and Orthogonal Multiplexing Techniques," *ALOHA System Technical Report B74-3*, University of Hawaii, June 1974.
- [30] STOUTENMYER, DAVID R.: "Analytical Optimization Using Computer Algebraic Manipulation," *ALOHA System Technical Report A74-3*, University of Hawaii, June 1974.
- [31] BINDER, RICHARD: "A Dynamic Packet-Switching System for Satellite Broadcast Channels," *ALOHA System Technical Report B74-5*, University of Hawaii, August 1974.
- [32] LIN, SHU and YIU, KAI-PING: "An Improvement to Multifold Euclidean Geometry Codes," *ALOHA System Technical Report A74-4*, University of Hawaii, August 1974.
- [33] BINDER, RICHARD: "The ALOHANET MENEHUNE - Version II," *ALOHA System Technical Report B74-6*, University of Hawaii, September 1974.
- [34] BINDER, RICHARD: "ALOHANET Protocols," *ALOHA System Technical Report B74-7*, University of Hawaii, September 1974.

- [35] WAX, DAVID: "Status Report on UH/ALOHA Participation in the ATS-1 Computer Communications Experiment," *ALOHA System Technical Report B74-8*, University of Hawaii, September 1974.

TECHNICAL NOTES

- [36] YAO, ANTHONY: "A Software Buffer-Control Unit for The ALOHA System," *ALOHA System Technical Note 72-1*, University of Hawaii, May 1972.
- [37] SUSMAN, HOWARD: "Plotz -- A Terminal Plotting Subroutine," *ALOHA System Technical Note 72-2*, University of Hawaii, August 1972.

MISCELLANEOUS

- [38] ABRAMSON, NORMAN: "ARPANET Satellite System," *ARPANET Satellite System Note 1, NIC Document 11283*, Stanford Research Institute, Stanford, California, 1972.
- [39] ABRAMSON, NORMAN: "ARPANET Satellite System," *ARPANET Satellite System Note 2, NIC Document 11284*, Stanford Research Institute, Stanford, California, March 1972.
- [40] ABRAMSON, NORMAN: "ARPANET Satellite System," *ARPANET Satellite System Note 3, NIC Document 11285*, Stanford Research Institute, Stanford, California, April 1972.
- [41] ABRAMSON, NORMAN: "ARPANET Satellite System," *ARPANET Satellite System Note 4, NIC Document 11286*, Stanford Research Institute, Stanford, California, May 1972.
- [42] ABRAMSON, NORMAN: "ARPANET Satellite System," *ARPANET Satellite System Note 5, NIC Document 11287*, Stanford Research Institute, Stanford, California, May 1972.
- [43] ABRAMSON, NORMAN: "ARPANET Satellite System," *ARPANET Satellite System Note 6, NIC Document 11288*, Stanford Research Institute, Stanford, California, May 1972.
- [44] ABRAMSON, NORMAN: "ARPANET Satellite System," *ARPANET Satellite System Note 7, NIC Document 11289*, Stanford Research Institute, Stanford, California, May 1972.
- [45] BINDER, RICHARD: "Effects of Retransmission Delay on the Degradation of an ALOHA Channel," *ARPANET Satellite System Note 22, NIC Document 12166*, Stanford Research Institute, Stanford, California, November 1972.
- [46] ABRAMSON, NORMAN: "Excess Capacity of a Slotted ALOHA Channel," *ARPANET Satellite System Note 26, NIC Document 12735*, Stanford Research Institute, Stanford, California, November 1972.

- [47] ABRAMSON, NORMAN: "Excess Capacity of a Slotted ALOHA Channel (continued)," *ARPANET Satellite System Note 30, NIC Document 13044*, Stanford Research Institute, Stanford, California, December 1972.
- [48] BINDER, RICHARD: "Another ALOHA Satellite Protocol," *ARPANET Satellite System Note 32, NIC Document 13147*, Stanford Research Institute, Stanford, California, December 1972.
- [49] WAX, DAVID W.: "The ALOHA Radio Modulation Scheme," *Packet Radio Temporary Note 17, NIC Document 13863*, Stanford Research Institute, Stanford, California, January 1973.
- [50] BINDER, RICHARD: "General Description of the University of Hawaii ALOHA System MENEHUNE," *Packet Radio Temporary Note 18, NIC Document 13877*, Stanford Research Institute, Stanford, California, January 1973.
- [51] KANEHIRA, EARL: "Transient and Steady-State Analysis of Reservation ALOHA," *ARPANET Satellite System Note 35, NIC Document 14180*, Stanford Research Institute, Stanford, California, February 1973.
- [52] LU, SHYUE-CHING: "Dynamic Analysis of Slotted ALOHA with Blocking," *ARPANET Satellite System Note 36, NIC Document 14790*, Stanford Research Institute, Stanford, California, 12 March 1973.
- [53] BINDER, RICHARD: "Acknowledgement Schemes in an ALOHA Channel," *Packet Radio Temporary Note 44, NIC Document 15687*, Stanford Research Institute, Stanford, California, April 1973.
- [54] ABRAMSON, NORMAN: "The Spatial Capacity of an ALOHA Channel," *Packet Radio Temporary Note 49, NIC Document 16087*, Stanford Research Institute, Stanford, California, April 1973.
- [55] WAX, DAVID W.: "ALOHA Impulse Noise Measurements," *Packet Radio Temporary Note 50, NIC Document 16725*, Stanford Research Institute, Stanford, California, May 1973.
- [56] MURTHY, PARIMI: "A Comparison of Spread Spectrum with ALOHA for Packet Radio Networks," *Packet Radio Temporary Note 58, NIC Document 16750*, Stanford Research Institute, Stanford, California, May 1973.
- [57] GAARDER, THOMAS N.: "Some Notes on ALOHA Channels with Capture," *Packet Radio Temporary Note 60, NIC Document 16753*, Stanford Research Institute, Stanford, California, June 1973.
- [58] CHEN, PO-HSIIN and WAX, DAVID W.: "Loss Measurement of Field Strength for ALOHA Channel," *Packet Radio Temporary Note 70, NIC Document 17894*, Stanford Research Institute, Stanford, California, July 1973.
- [59] CHEN, PO-HSIIN and WAX, DAVID W.: "ALOHA Noise and Propagation Loss Measurements," *Packet Radio Temporary Note 78, NIC Document 19702*, Stanford Research Institute, Stanford, California, November 1973.

- [60] WAX, DAVID W.: "Status Report on UH/ALOHA Participation in the ATS-1 Computer Communications Experiment," *ARPANET Satellite System Note 56*, *NIC Document 20805*, Stanford Research Institute, Stanford, California, 28 February 1974.
- [61] KUO, FRANKLIN F.: "INWG Workshop Report: Summary of an Ad Hoc Committee Report on Political, Social and Economic Issues on Network Interconnections," *International Network Working Group Note 52*, *NIC Document 22030*, Stanford Research Institute, Stanford, California, 19 March 1974.
- [62] MURTHY, PARIMI: "Equalizers for PRN," *Packet Radio Temporary Note 79*, *NIC Document 19716*, Stanford Research Institute, Stanford, California, May 1974.
- [63] OKINAKA, ALAN and WAX, DAVID W.: "Preliminary Design for an ALOHA Repeater," *Packet Radio Temporary Note 101*, *NIC Document 30486*, Stanford Research Institute, Stanford, California, May 1974.

5.0 PERSONNEL

Faculty

Norman Abramson
Rahul Chattergy
Franklin F. Kuo
Shu Lin
Vincent J. Nicholson
David Pager
Edward J. Weldon

Technical Staff

Karen Ah Mai
Richard Binder
Michael J. Ferguson
Christopher G. Harrison
Gary T. Mikasa
Peter L. Nelson
Alan K. Okinaka
Howard D. Susman
David W. Wax
Thomas C. Westphal

Casual

Robert J. Reininger

Graduate Assistants

Wilfred S. Ageno (also student help)
Moon-Suk Ahn
Richard H. Case (also student help)
Yoshihiko Ebihara (also casual)
Dieu Huynh (also student help)
Earl M. Kanehira
Prem N. Kapoor
Lynn M. Kennedy
Wai Sum Lai (also student help)
Peter S.W. Law
Wrenwick K. Lee
Parimi S.B. Murthy (also student help)
Lars E. Ottosson (also student help)
Lorraine T. Sakaguchi (also student help)

Student Help - Technical

Lee E. Castonguay
Po Hsiin Chen
Theodore Y.H. Chinn
Gary W.C. Chun
Richard H.C. Chun
Sushil K. Garg
Steve Glanstein
Jeffery H. Hara (also casual)
Neil Y. Iwamoto
Branden B. Johnson
Keith K.L. Luke
Alan H. Maruyama
Marsha H. Murai
Clifford K. Nishimura
Rodney E.K. Okano
Richard T. Schiavoni
Russell P. Seeney
Don J. Slepian
Izumi Soma
Dennis A. Stone
Alan W. Sugahara
Amy E. Tanaka
William V. Tift (also casual)
James M. Ulyshen
Morris H. Wilson, Jr.
Anthony C. Yao
Wendy N. Yin

Administrative Staff

Donald L. Dobbs
Jared S. Ikeda
Margaret M. Iwamoto
Evelyn G. Kwock (also student help, casual)
Doreen T. Leland
Jeanne H. Moriyasu

Administrative Student Help

Alva A. Chai
Noreen P. Ettl
Sharen M. Gomoto
Laurie L. Kino
Lily F. Ko
Susan Y. Koga
Nancy Y. Nakagawa
Liana Petranek
Beryl E. Tanaka
Carole A. Utsumi
JoAnn C.W. Wong

II. TASK II

1.0 INTRODUCTION

This section covers the activities of Task II during the period from November 1, 1971 through October 1, 1974. Task II, Research in Multi-processor Computing Systems, was formed under Contract NAS2-6700 and added to THE ALOHA SYSTEM in late 1971 to investigate multiprocessing system techniques and make possible the transfer of technology incorporated in the BCC 500 computing system to the general computing community and to specific ARPA-supported endeavors.

As circuit speeds become limited by that of light and by packaging considerations, some form of concurrency in processing is the only alternative for cost/performance improvements. The large "super computers" rely on the programmer's ability to organize his data such that a single machine instruction can direct a number of arithmetic units to do similar computations simultaneously (as in ILLIAC IV), or on the engineer's ability to design hardware to execute a number of instructions simultaneously (as in STAR or the TI machine). Many jobs do not lend themselves to these approaches, however, especially those which interact with persons at remote terminals. Systems having more than one processor are not new to the computer scene, nor are those running independent jobs directed from interactive terminals (time-sharing). Most such systems use the additional processors either in the identical role of the main processor or to do system I/O. The system of Berkeley Computer Corporation was the first serious attempt to distribute the functions of system management (the operating system) over several simpler, dedicated processors in the rigorous environment of interactive computing.

The BCC 500 system was designed at Berkeley Computer Corporation in 1969 using principles and philosophies of multiprocessor systems developed by researchers at Project GENIE of the University of California at Berkeley. A number of these individuals helped form the company upon leaving the Project. In the new environment all of the development work--detailed design, fabrication and testing--was done both in hardware and software areas, and the Project technology thus moved from the realm of theory to that of practice. BCC was forced, however, by financial difficulties to terminate operations in early 1971 after constructing a working prototype. Thus the opportunity existed to bring to Hawaii all of BCC's equipment and designs and a number of former BCC personnel to set up and operate the system, publish aspects of it, consult with ARPA contractors having multiprocessor systems, and establish a basis for doing further research into multiprocessor systems. A proposal to this effect was submitted in mid-1971 for a three-year effort beginning in September. The proposal was accepted, although a two-year contract (NAS2-6700) was given.

During the period of proposal evaluation and contract negotiations, however, a slippage in completion of new facilities to house the Task occurred, delaying the Task's effective starting date some five months. In February 1972 the equipment was moved from Berkeley to Honolulu and placed in the new Holmes Hall, still under construction. Work began on the system's installation and refurbishment (it had deteriorated somewhat during its period of storage) in March. Software work began in May. In February 1973 the system first ran using a temporary CPU. This permitted software activities to enter the checkout phase.

From February through August 1973 (the end of the two-year period), the hardware was improved and much software work was accomplished. The work was performed by the Task II staff most capably and under adverse conditions (initially there was no plumbing, telephone, air conditioning, or elevator in the building). Since the staff was acquired more locally than had been originally planned, the new people required training. Thus, at the end of August 1973 the work on the installation was not complete.

A proposal was submitted to continue the work of the Task in the same vein of the earlier proposal: to permit a few month's more effort on the 500 completion and documentation and then to enter the research phase. NAS2-6700 was continued for another year; the Task was asked by ARPA to complete its work on the 500 system in four months, but resubmit its proposal for new directions in research.

The topic chosen was operating system security; specifically, methods for securing the TENEX-based operating system proposed by the Institute of Advanced Computation (IAC) at NASA/Ames Research Center. This system was and is still under development. It includes the ILLIAC IV computer and the UNICON trillion-bit store, two large and expensive resources which are being used from different locations in the United States by means of the ARPA network. There is a desire by DOD that in the near future at least the ILLIAC be made available for the processing of classified information. Our proposal to study various means for providing the requisite security via TENEX was accepted and support was provided for the remaining eight months of the contract year.

Section II will be only a brief summary of the work accomplished by the Task. Each subsection includes references to technical documentation which has been produced under the Contract and which has been supplied to ARPA and IAC.

2.0 BCC 500 WORK

This work was accomplished during the first 26 months of the Task. It may be described in two, separate major activities: hardware work and software development. These will be discussed separately. The two efforts were closely interlinked, of course, and influenced each other, sometimes detrimentally.

2.1 Hardware Work

The prototype BCC 500 system was built in Berkeley, California on the assumption that the system would never be moved. It was built into open frames in two different rooms. To move the system to Honolulu required total disassembly, and this necessarily caused considerable damage. Some sections could be disassembled only by cutting and destroying cables.

Several man-months were spent in arranging for the move and in packing the equipment. The equipment, in crates, filled a chartered 707 jet freighter. Less than \$100 of damage was done to the equipment in the complete move--a tribute to the persons responsible for planning of the move and packing.

The BCC 500 system in Berkeley consisted of a total of nine processors, six of which were part of the central system. The others were intended to be connected to the central portion by means of communications interfaces. The configuration chosen for Honolulu did not require the remote processors, however. Thus, only the central portion of the system was made operable.

Also connected to the central memory were two drum units and a dual-positioner disk file. Non-terminal I/O was handled by a 360/30 with its attendant peripherals. The 360 was not part of the BCC equipment and had been returned to IBM earlier.

Two of the processors (CPUs) were dedicated to execution of user processes. One of these had not yet been completed; its role was performed by a temporary CPU called the SMP (system measurement processor), destined later to be used for system utility purposes. The other processors were dedicated to running aspects of the overall operating system. The MSCH (microscheduler) primarily scheduled the CPUs and handled process wakeups. The CHIO (character input/output) performed input/output activities for the terminals. It was intended to communicate with the remote processors. The AMC (auxiliary memory control) did all of the system memory management.

The processors communicated through common areas on the central memory using a simple hardware interlock mechanism called PROTECTS. Each processor was implemented around a basic microprocessor structure capable of performing operations at peak rates of around 15 million per second. The operating system processors were guided by algorithms written directly into microcode (a unique aspect of the system) such that their accesses to memory were for data only. In effect, each was (and is) a special-purpose, data-driven processor capable of very high speed operation (as compared with a processor running conventional software).

The CPUs were implemented on similar microprocessors, the microcode in this case emulating a rather elegant processor designed in conjunction with a higher-level language aimed at systems programming (SPL). No assembler was produced for this CPU; code for it was to be written only in SPL and other languages as they became available. The CPU was equipped with features to make it efficient when running SPL and FORTRAN. It was also equipped with a mode which emulated the user mode of the XDS 940 system. Thus, all 940

software could be run on the system. (A software package called the "940 Emulator" fielded the system calls and converted them into appropriate 500 system calls.)

All of the above-mentioned hardware was received in Honolulu physically intact, but far from operable. As each hardware item was undergoing renovation it was scrutinized for design flaws. A number of design weaknesses were discovered and corrected. In many cases this involved adding wires to existing printed circuit boards; in some cases, however, new boards were required altogether. The items were then checked out as completely new devices and installed into new cabinets designed to provide better cooling.

The following hardware projects were accomplished during the contract:

2.1.1 Site Preparation

A number of projects were required to prepare the space to receive and house the equipment.

Room Layout

The space available for the equipment was not overly large and was rather oddly shaped. After some difficulty, the equipment was appropriately laid out into the room in such a way that it could be worked on and the room could be properly cooled.

Cabinets

A completely new set of cabinets were acquired to specifications developed specifically for the machine. These cabinets were arranged in an H-structure with the processors and other computer electronics in the front sections of racks and the power supplies and core memory in the back section. The two were connected by a narrow cable raceway.

AC Power

It was necessary to specify and pursue the installation of new power facilities in the building to provide for the operation of the equipment. Lengthy specifications for this power were developed and put through the State bidding process and final installation.

Motor Generator

The BCC equipment included a 37.5 KVA motor generator with a large fly-wheel for ride-through operation. The motor generator removes switching transients from the primary power and provides excellent voltage regulation. The fly-wheel provides enough inertia to operate the generator for about 15 seconds after a primary failure. The primary power can fail for up to 5 seconds without forcing a shut-down. Power from the unit was carried to a distribution panel in the machine room to feed the various system electronics components (the drum and disk units were powered directly from primary power).

DC Power

Each identifiable unit in the system was provided with its own set of power supplies. These supplies were monitored and switched on and off by a set of power sensing and control units. The original BCC power units were altered to provide improved reliability and were renovated before installation in the new cabinets.

Air Conditioning

A 15-ton fan coil unit was installed in the room to provide cooling for the system. The unit was connected to the building chilled water system. This was done primarily for reasons of economy. Experience since then has shown that the building system is not well regulated and is subject to failure, especially when there are power problems. This has affected the system reliability.

Safety Interlocks

Experience with equipment of this nature has shown that its reliability is greatly impaired when it is switched on and off frequently. To operate the equipment in the absence of personnel, however, without requisite interlocks is clearly dangerous. Thus, the room was equipped with a number of safety circuits designed to totally shut down the room automatically in the event of air conditioning failure, or a power failure, and manually by pressing a single button located at each entry to the room.

2.1.2 Hardware Renovation

Each of the major units of the BCC 500 system was renovated as previously discussed.

Microscheduler

The microscheduler was initially set up in a test arrangement and served as the model for analysis of the basic microprocessor design. In this way, a number of timing bottlenecks and logic design flaws were discovered and corrected. The microscheduler was then completely refurbished and installed in the cabinets. The processor was provided with ROM parity and a number of other new features.

CHIO

In the same manner the CHIO was refurbished and installed in the cabinets. The microcode was modified to permit concurrent operation of the CHIO and the common test processor (CTP), a portion of microcode emulating the instruction set of a processor similar to the Xerox 940. This capability permits software to be executed on the CHIO on a time-shared basis with its other activities to facilitate handling the ARPA network

and/or other input/output not originally provided for in the system. The software may communicate with the original microcode by means of special calls.

AMC

This processor was similarly refurbished, installed in the system, provided with ROM parity, etc.

AMTU

This unit contains all the drum/disk transfer and fast memory interface electronics. The unit was deemed to be in adequate condition and was installed and checked out without refurbishment. Replacement of the unit was deemed impracticable because of its uniqueness in the system and its complexity. The unit has since operated flawlessly.

Drum Ø

This Bryant 1100-head drum unit was received as a newly refurbished device from its manufacturer. It was placed into operation early in December 1972 by the manufacturer's installer. Shortly thereafter it suffered a number of head crashes involving about 7% of the available storage area. Since the unit was not under warranty, it was necessary for Task II personnel to remedy the situation. Investigation showed that the crashes had occurred because of dirt within the drum. Because of very dirty conditions in the room, and of several possible means for entry of dirt into the unit, it seemed that the drum would require complete cleaning every two months if it were to survive. Moreover, our previous experience with such a drum was that after such a crash the drum had, at most, a few months of life remaining.

A number of measures were developed and taken to forestall future accidents. The machine room was thoroughly cleaned and stern measures to keep it that way were instituted. Carpeting was installed on the floor and all personnel were required to remove their shoes upon entering the room, thereby greatly reducing the amount of dirt brought into the room and eliminating problems with static electricity. A built-in vacuum cleaning system was installed so that dirt removed from the carpet was not transferred into the room atmosphere. The room was pressurized against dirt entry by air conditioning modifications. The drum was sealed more carefully and was also pressurized by dried air passed through a micron filter. The method of cooling the unit was changed to eliminate the fans that were thought to be the major source for dirt entry.

The ruined heads were removed and a number of new heads installed in spare slots. That, together with existing spares, made up for the loss of storage and the drum was restored to full capacity. An aerosol particle detector was fitted to the unit. This proved invaluable, not only in serving as a sentinel during operation of the unit, but also as a means of locating heads not flying properly. All of this effort has worked well. The unit has functioned since without incident and has not required servicing of any kind.

The measures described in renovating and protecting this drum have been transported to the ILLIAC IV NASA/Ames IAC TENEX system, which includes a drum of similar manufacture. This drum, operating in a similarly dirty environment, has had even more serious maintenance problems.

Drum 1

This unit was received originally in the equipment from BCC where it had run acceptably for months. It suffered the same fate as drum 0, however, losing approximately the same number of heads. All of the measures described above for drum 0 were applied to drum 1 and the unit was made operational.

Drum Electronics

Both drum units are 24-bit parallel devices and there is considerable electronics associated therewith. This equipment was in acceptable condition, requiring only cleaning and checking; minor modification was made on a few of the printed circuit boards during check-out to improve signal quality.

Disk File

The Bryant disk file was completely rebuilt on site in Honolulu by Bryant personnel. A complete set of "crash proof" disk surfaces was installed, and the original disks were stored against future need.

During this time one of the two actuators failed and replacement was required. The disk and its associated electronics began to perform acceptably with relatively little attention. Subsequently, this performance has been improved by systematic adjustment of various heads on the file. The file has been operated for approximately two years without incident except for one head crash which occurred due to excessive dirt. The head was destroyed in the crash, but the surface was found to be still useful. Consequently, the head was merely replaced and the file was restored to operation. The room cleanliness measures previously described have forestalled a recurrence of the situation.

Water Chiller

The disk file required a 5-ton water chiller for environmental control. This water chiller was installed in a utility area on a floor above at the time the disk was installed.

MPMBM (Microprocessor Memory Bus Multiplexer)

The MPMBM was deemed to be in such poor condition that it was not worth renovation. Consequently, a completely new unit was designed, constructed, and checked out.

Fast Memory

The fast memory is one of the more complex and critical portions of the system. It consists of four "quadrants" of electronics connected to a total of eight Ampex core memory modules and serves to buffer and schedule memory requests from up to four conflicting sources over the eight core modules. It contains 32 registers of active storage, providing a modest look-behind capability. Various printed circuit boards of the unit were refurbished and a number of design faults were noted and corrected. This resulted in the necessity for reconstruction of a small number of printed circuit boards. The means by which the original fast memory was connected to the various processors (i.e., the memory cabling) was considerably damaged during the disassembly. Because new cabling was required and because of the doubtful quality of the fast memory back plane it was decided to replace the back plane. This was done utilizing programs which ran on the system itself with the original fast memory configuration connected temporarily. The system was then idled for about six weeks while the memory was reconfigured on the new back plane and with new cabling and connectors.

Core Memory

The Task received 11 Ampex R-G core memory modules, each of 16 K by 50-bit capacity, one microsecond cycle time. Several of these modules were not functional, as some deterioration of the electronics components had occurred. One had not yet been equipped with extensive modifications designed and installed by BCC on the others. This module was updated and all the modules were placed in working condition.

An off-line testing facility for Ampex modules was designed and fabricated using a discarded BCC microprocessor prototype back panel and old printed circuit boards. This facility, equipped with a teletype interface, permits modules to be operated in test mode in a variety of ways while off-line, thus assuring a functional replacement should one of the active modules fail.

CPU 1.5A

This unit was refurbished and installed in the cabinets. It gave us a 40% increase in computing power over the temporary CPU when it went into operation. The unit is implemented on a microprocessor except that it is provided with a hardware multiplier, instruction fetch unit, and a 128-entry memory map. The processor was also equipped with ROM parity. As this was the first processor so equipped, it was necessary to produce, update, and fully verify the processor's microcode source language to include correct parity. These procedures were then used to produce updated and verified microcode for all other processors, assuring fully documented microcode for all processors.

CPU 1.5B

The unit was refurbished and installed into the cabinets. This unit had never been completed by BCC and required slightly more work than its companion unit. It was the last processor to be placed in the system on-line and effectively doubled the system's computing capacity.

2.1.3 System Integration

Central Control

This unit was redesigned and built to replace an older substandard BCC component. It contains clock generation and distribution and a number of circuits used in conjunction with the maintenance panel.

Real Time Clock and Battery Charger

The system real time clock was refurbished and installed in the cabinets. An accompanying power supply unit including a nickel cadmium battery and charger was refurbished and installed. The battery serves to provide a continuous source of power for the system real time clock and a few other critical circuits.

Console

An operating console consisting of switches, lights, and push-buttons connected by a long cable to the system central control was designed and built to replace an obsolete unit.

Maintenance Panel

A special maintenance panel was built with a number of lights to indicate various error or abnormal conditions within the system and a number of switches to permit various actions related to hardware maintenance and system configuration.

Memory Cabling

The original system memory cabling which tied together all of the processors and drum disk controllers to the fast memory, and the fast memory to the core modules, was virtually destroyed as the system was disassembled. In the process of reconstruction of the fast memory, a new set of cables was built. These cables were built to different specifications than the original cabling, using a special form of flat cable consisting of alternate ground and signal lines. Paddle boards were also designed to accommodate the termination and connection of these cables to the printed circuit boards of various units. We were pleased to observe only a minor amount of cross talk on the cables and an extremely high quality signal. Previously 100 ohm coax was used for each signal line. This kind of cabling was extremely bulky and quite difficult to terminate for connection purposes.

CHIO Multiplexer

A completely new CHIO multiplexer was fabricated from components which had been originally destined for the Phase II CHIO multiplexer to be used in connecting the CHIO processor with the remote DCCs. This multiplexer was thus designated the CHIO multiplexer 1.5 and served to replace the Phase I multiplexer which was of substandard quality. The unit consists of a back panel accommodating printed circuit boards for various kinds of communication line connections to the CHIO. The CHIO is currently equipped with 16 local terminal (current-loop) connections and eight RS 232 terminal connections. It is also equipped with an interface to the HP2100, to a PDP-11 (via a 9600 bps RS 232 modem interface), and to the ARPA network IMP (via a specially designed 50 Kb interface).

PROTECTS

Analysis of the operation of the system revealed that the PROTECTS, a system of processor interlocks to permit the protection of system objects like global tables, were inadequate as originally implemented by BCC. This necessitated redesign of the entire PROTECTS system. The new design involved changes in the processors as well as the production of four printed circuit boards in the system central control area. The new PROTECTS include a self-checking feature, providing a warning indication in the event of a failure, and feature an equal priority circuit giving each processor equal treatment on the average in contention resolution.

HP 2100A and Potter Tape Unit

This commercial minicomputer was received and interfaced to the CHIO multiplexer. A controller for the accompanying tape unit was developed by a commercial supplier to Task specifications. The tape unit, a 120 ips, 800 bpi, 9-track unit is used primarily as a back-up facility for system files and for communication with other systems. It was by this means that the BCC software, brought to Hawaii on tape, was reintroduced into the system.

Iomec Line Printer

An Iomec 200 lpm line printer was received and interfaced to the HP2100A. This interface was designed and built as two sub-units, one at either end of a 100-foot cable linking the units. The printer is maintained in an adjoining area primarily to keep people away from the immediate area of the system and to prevent paper dust from interfering with drum and disk operation.

2.1.4 Documentation

The state of hardware documentation is good.* Accurate logic diagrams exist for all hardware. There are correctly marked Copies of Record which are kept in the machine room with the equipment, and back-up originals stored in a separate area. During the contract all microcode was verified and reconciled to old Copies of Record. During this effort several existing microcode errors were located and corrected. The microcode source language for every processor was updated and recompiled. Finally, the recompiled microcode was checked against the Copies of Record and against the actual microcode boards themselves. During the contract a wirelisting program written for the Xerox 940 was obtained from Xerox PARC and modified slightly to be used with BCC style back panels. This program was then used to generate up-to-date wirelists for each of the back panels in the system. A number of other hardware documents were produced describing the hardware and various indicators and protection devices installed during the period of system integration. There also exists stated maintenance procedures to be followed for the various units of the system.

2.2 Software Development

The software accompanying the hardware was stored in symbolic and binary form in BCC 500 system files on a number of 7-track 556 bpi magnetic tapes. For compatibility with the University Computing Center machines, it

* No hardware documentation is included with list of Task II documentation at the end of Section II, as it is considered overly specialized. The documentation consists of logic diagrams, microcode listings, and detailed hardware operation descriptions.

was determined that the BCC 500 tape unit should be an 800 bpi, 9-track unit. Thus, it was necessary to copy all of these tapes from one format to the other. A number of tape handling programs were written for this purpose. During this same period (i.e., before the hardware was available) programs for the IBM 360 were written for the purpose of reformatting some of these files and listing them. This gave the programmers access to the BCC software before the hardware was available. The original BCC 500 configuration included a direct connection to an IBM 360 Model 30 which operated a number of tape units, line printers, card readers, etc. To replace this arrangement, the HP 2100 tape unit and line printer were acquired. It was necessary to write and debug software for the 2100 to operate as a controller for all of the devices, driven by requests originating in the 500. This effort was done so as to make the 2100 emulate precisely the past behavior of the Model 30, thus requiring no changes in BCC software.

The state of the operating system received from BCC was relatively good. The system was interim and heavily patched, but its bugs had been mostly located and the patches were documented. Thus, after a modest effort to find portions of the software scattered on various tapes, collect them together, and remake the various patches, the system was ready to run well. As the hardware improved, so did the reliability of the system. A few new software errors were discovered and corrected.

The primary software accomplishment was the production of a working SPL compiling system. The compiler originally used by BCC was written to run on the Xerox 940 and, while also runnable on the BCC 500, ran with great inefficiency. This compiler had been written at BCC in QSPL, a 940 language.

An SPL written in its own language had been partially completed, but had never been debugged or documented. Production of the SPL first required checking and full documenting of about 400 pages of listing. A number of errors were found and corrected in this work. Then compiling and on-line debugging began. As a partial check on the correct operation of the compiler, it was demonstrated that it would correctly compile itself. The SPL was then used to recompile and modify a number of BCC 500 software packages. The final version of SPL contains about 20,000 source language statements.

The system utility was enhanced to provide a number of new features such as terminal linking. This required concurrent microcode changes, as the original microcode was in error. Other utility subsystems were recompiled and updated, such as a subsystem permitting use of the magnetic tape and implementing a stand-alone file system for the disk file which is almost impervious to system crashes. A number of 940 subsystems were also operational at the end of the year; these required no effort since the CPU has a 940 mode and executes 940 user code directly. These subsystems include a comprehensive text editor, QED; NARP, a macro assembler for the 940; DDT, a loader-debugger; an interactive SNOBOL IV; a JOSS-like interactive programming language called CAL; a RUNOFF[10], and others.

2.3 Other Activities

The Task worked on a number of matters which were related to our research charter but are best described under a miscellaneous category.

2.3.1 Connection of the BCC 500 System to the ARPA Network

Originally it was planned to connect the system to the ARPA network via an NCP running on the CHIO processor. Work in better documenting the NCP specifications progressed to the point at which it was possible to write an NCP in microcode. This was done and it was noted that the additional code could not fit into available space in the CHIO. Another NCP was written in QSPL - a 940 systems programming language - and a near-working version of a CHIO-based NCP was brought up in a short time. This version ran on the CTP (940 emulation) facility, part of the CHIO.

At this time the Task was asked by IAC to serve as a test-bed for the ELF-based NCP under development at Ames. ELF, a system designed at UC, Santa Barbara for the PDP-11, is equipped with an NCP and TELNET; it connects with the "host" via synchronous and asynchronous communication lines. Unfortunately, the ALOHA-TIP was equipped only with two external IMP ports, and one of them was being utilized by the radio-terminal ALOHA system. It was decided in view of the greater potential for utility of ELF that the CHIO-NCP would be indefinitely postponed. As the contract ended, the ELF hardware including interfaces both with the TIP and the BCC 500 was fully operational; ELF software at Ames was nearing completion.

The 500 has nevertheless been utilized via ARPANET since its early operation by means of special TIP lines made to operate in a transparent mode. Both local and mainland users have regularly used the 500 system via the TIP in this fashion.

2.3.2 Local Echoing and Terminal Management in a Satellite Environment

The designers of the BCC 500 were among those to adopt and stress full duplex, character-by-character communication between terminal and computer. This approach is, of course, now used by a number of systems on the ARPA network, but as TIPs, IMPs, and local and remote hosts get involved in the scheme, such means of communication appears less attractive because of the cumulative delay for the echos. The same phenomenon appears when a long time-delay transmission line is used--a satellite link, for example, where the round-trip time approaches 500 milliseconds.

The BCC 500 was designed to have small terminal-handling computers located physically near groups of terminals (where the transmission delay is negligible). These devices produce the echos for full duplex terminals using certain strategies, so that the transmission time to the 500 is much less a factor. This situation is not unlike the ARPA network in which the TIP or local host can play the role of the local echo generator and the remote host that of the main machine.

During the contract, the echoing strategies in the 500 system were reviewed for possible use on the ARPA network and over satellite links. This work resulted in a scheme which was documented, and submitted to the Network Working Group [12], and in a later technical report [6].

2.3.3 Participation in COTCO Study

The Task was asked in July 1973 by NASA/Ames Institute for Advanced Computation to do a study on the proposed COTCO experiment. Considerable effort was expended over a short period, and a draft report was submitted to IAC in August of that same year [31].

2.3.4 New Microprocessor Design

In response to needs felt by NASA/Ames IAC, a pilot study was launched to determine the feasibility of redesigning (more accurately, repackaging) the BCC standard microprocessor for use in the ILLIAC IV system. The IAC system design called for the use of a number of small processors to perform system management functions similar to the BCC 500's CHIO, AMC, and MSCH functions. The PDP-11 was chosen for this application, but as development progressed it became apparent that the PDP-11 was not sufficiently powerful. The notion was to borrow technology used on the 500; to use the processor microcode directly for implementation of frequently-called, fixed functions and to use a standard emulated CPU instruction set to run classical software, more readily changed. The addition of a function call instruction allows the processor to escape from emulation to direct microcode execution of the otherwise high overhead functions. The main point was that by enhancing a, say PDP-11, with direct microcode capability of the class of the 500's microprocessor, its power could be greatly increased while maintaining its basic character as a PDP-11.

The study required about six man-months of effort and resulted in the conclusion that such a processor could be produced easily using available packaging schemes and MSI logic. Six additional man-months produced by the end of the contract a complete set of logic diagrams for the processor and a section of logic called RAID (Remote Assistance in Debugging) which permits another PDP-11 to be connected to the processor via its UNIBUS and to perform complete hardware diagnostics. The RAID capability is of novel design and looks promising as a means for finding almost any type of processor failure.

A modest software support effort was carried out in conjunction with the processor design. A wirelist program was developed for IAC from existing BCC software, modified to accommodate the Augat-type boards chosen for the processor. We were assisted in this greatly by IAC personnel (more accurately, it was a cooperative venture). Also beginning with techniques developed by BCC, a processor microcode compiler-simulator was developed for use by IAC. The higher-level language for the microcode had been developed at BCC. This language was only slightly modified, and the compiler-simulator was written in SPL. It allows microcode, after compilation, to be simulated.

2.3.5 Remote FTP

To provide a means of transporting files to and from the 500 system by means of the ELF-NCP, a specification was written for an FTP capability suitably distributed between the 500 system and ELF [11]. At the end of the contract the 500 code had been written and debugged, and we were awaiting the ELF code from IAC. This effort, if concluded as we hope, may prove useful to other sites attempting connection to the ARPA network. A paper describing the approach has been submitted to the upcoming NCC (1975)*.

2.4 System Documentation

From September through December 1973, Task II project concentrated primarily on the production of adequate system documentation. Some documents were brought from BCC and required only minor modifications. Other documentation was totally lacking. A number of technical reports and manuals were produced in this effort [1,2,3,7,8,9].

* Paper written by J. McConnell of IAC and D. Yonamine of Task II.

A full session was devoted to the system in the 1974 Hawaii International Conference on System Sciences. A total of six presentations describing the system organization and operating system were given in the session.

2.5 System Operation

The BCC 500 system first functioned as a full system in February 1973. Initially the system was useful only for system programming purposes, as reliability was poor. In March 1973 the system was placed on a schedule of four user hours/day, the remainder of the time devoted to hardware activities or systems development activities. The schedule for guaranteed user time was subsequently modified on several occasions, but continuously expanded so that by May 1974 the system was operated 24 hours/day, seven days/week, available to users continuously except for Saturday, reserved for hardware and software maintenance purposes.

System reliability was at the level of better than 95% up-time during those hours when Task personnel were present to assist with system operations. Even though the system functions a great deal of the time unattended, it has nevertheless provided overall up-time of approximately 80%.

3.0 OPERATING SYSTEMS SECURITY WORK

In January 1974 Task II began work under its new charter in support of the ILLIAC IV computer system at the NASA/Ames Institute for Advanced Computation. The work was directed mainly to a thorough analysis of the TENEX operating system with a view toward reorganization of most of the code (and replacement of some) so that the TENEX virtual machine could be integrated into the planned I4 operating system with improved security features.

The major conclusions and results of these investigations were:

- Securing TENEX in the sense of providing certifiable multi-level security within it should be abandoned as requiring far too much effort (perhaps three years for a project our size).
- Separating the TENEX monitor into critical and non-critical parts and isolating each in a separate hardware "ring" was crucial to provision of multi-level security, but was not cost-effective enough to justify pursuing it for its own sake.
- Removal of certain management modules from the monitor into management processors would contribute enough to system efficiency and reliability that it should be seriously considered for incorporation in the IAC system.

The decision to abandon the idea of making a "secure TENEX" was reinforced by our observations of the difficulties being experienced by other projects which were attempting to retrofit security into existing operating systems, even in cases where a better hardware and system organization base existed than that of TENEX (e.g., MULTICS).

While this work was in progress it was brought to our attention that a specific need for secure operation of ILLIAC IV was being felt by the Navy's Acoustic Research Center (ARC) at Moffett Field, California. This group needed to process data of a classified nature on ILLIAC, but existing techniques for such situations would have required making the entire complex of machines--including TENEX--unavailable to users for many hours.

Accordingly, a straightforward method for providing for concurrent, but separate, operation of the two facilities was developed. The characteristics of this "encapsulation" plan, described in our proposal [PR AS 74010] as revised on August 2, 1974, were:

- It was limited in scope in that it would provide for secure use of the I4 only from or through the ARC facility.
- It was independent, as far as security goes, of the internal security and/or correctness and stage of development of IAC's operating system.
- It was comparatively simple. The amount of real research involved was small and therefore the risk of failure was comparatively small.

The work involved was pretty well charted at that time in our suggested Work Statement.

In addition to the above cited results, considerable documentation was produced and provided to IAC and ARPA [4,5,14-30].

4.0 DOCUMENTATION FOR PERIOD NOVEMBER 1, 1971 TO OCTOBER 11, 1974

TECHNICAL REPORTS

- [1] Wall, Charles F., Design Features of the BCC 500 CPU, R-1, January 3, 1974.
- [2] Lee, Wrenwick K., The Memory Management Function in a Multiprocessor Computer System--A Description of the BCC 500, R-2, September 12, 1974.
- [3] Freeman, Jack and John Davidson, BCC 500 Protection Mechanisms, R-3, May 15, 1974.
- [4] Lee, Wrenwick K., A Scheduling Processor for TENEX, R-4, September 4, 1974.
- [5] Wall, Charles F., Enhancements to TENEX protection Facilities, R-5, October 8, 1974.
- [6] Davidson, John, Partitioned Echoing in a System with Long Transmission Delays, R-6, November 25, 1973.
- [7] Lichtenberger, W. Wayne, and Jack Freeman, An Architectural Overview of the BCC 500 System, R-7, October 10, 1974.

MANUALS

- [8] Wall, Charles F. and Jack Freeman, BCC 500 CPU Reference Manual, M-1, December 31, 1973.
- [9] Wall, Charles F. and Judy Simon, BCC 500 SPL Language Reference Manual, M-2, December 31, 1973.
- [10] Lichtenberger, W. Wayne, PFS Reference Manual, M-3, May 1, 1974.

SPECIFICATIONS

- [11] Lichtenberger, W. Wayne, Specification for the Implementation of ARPA Network File Transfer Protocol by Means of the ELF System, S-1, March 15, 1974.

MISCELLANEOUS

- [12] Davidson, John, "An Echoing Strategy for Satellite Links," NIC Document 10599, Stanford Research Institute, Stanford, California, June 26, 1972.
- [13] Heckel, Paul C. and Butler W. Lampson, The BCC Terminal System, P-1, January 8, 1974.

TECHNICAL MEMORANDA

- [14] Kam, Alan, "Overview of TENEX File System Implementation, TM.74-29, September 26, 1974.
- [15] Kam, Alan, "Functions of TENEX JobØ," TM.74-30, September 25, 1974.
- [16] Wun, Enoch, "Basic Description of the TENEX Disk Structures, TM.74-27, June 21, 1974.
- [17] Wun, Enoch, "Disk I/O Request Service," TM.74-28, July 16, 1974.
- [18] Lee, Wrenwick K., "TENEX Terminal I/O," TM.74-31, September 4, 1974.
- [19] Simon, Judy, "TENEX Memory Management Code--Flowcharts," TM.74-32, October 3, 1974.
- [20] Jubinski, Paul, "The Balance Set Scheduling Function in TENEX," TM.74-33, September 10, 1974.
- [21] Simon, Judy, "Page Not in Core Trap," TM.74-17, April 8, 1974.
- [22] Freeman, Jack, "Notes on Compatibility and Critical/Non-critical Code Separation in the TENEX Monitor," TM.74-13, March 20, 1974.
- [23] Hironaka, Joan, Kam, Lynette and Charles F. Wall, "Documentation of TENEX Fork System Code," TM.74-34, September 30, 1974.
- [24] Wall, Charles F., "TENEX Protection Memorandum," TM.74-20, June 14, 1974.
- [25] Freeman, Jack, "Overview of the Implementation," TM.74-21, June 14, 1974.
- [26] Kam, Alan, "A Possible Implementation of Access Control Lists," TM.74-22, June 14, 1974.
- [27] Freeman, Jack, "JFN Protection," TM.74-23, June 12, 1974.
- [28] Freeman, Jack, "Management and Protection of Directory Numbers," TM.74-24, June 14, 1974.

- [29] Wall, Charles F., "Fork Protection Memorandum," TM.74-25, June 14, 1974.
- [30] Freeman, Jack, "Protected Programs: A Simple Implementation," TM.74-26, June 14, 1974.
- [31] Task II staff, "A Draft Report on Message Handling," August 1973.

5.0 PERSONNEL

Faculty

W. Wayne Lichtenberger

Technical Staff

Programmers

John M. Davidson (also grad asst, casual)
Jack Freeman
Alan C. H. Kam
Wrenwick K. Lee
Charles F. Wall

Engineers

Roger P. Bissonnette
Perr Cardestam
Donald E. Dodge
Lloyd T. Ebisu (also student help, casual)
Allen B. Goodrich
Rodney M. Inefuku

Graduate Assistants

Stephen Kasperowicz
Yu Hao Lin
Charles I. Shapiro
Judith H. Simon
Enoch W. K. Wun (also student help)

Operations & Technical

John R. Johnston
Peter L. Nelson
Dorothy R. Smith
Michael D. Waters
Thomas C. Westphal

Casual

Peter Hollister
James L. Ventura

Student Help - Technical

Terrence K. J. Ching
Leann R. Foss (also casual)
Steve Glanstein
Myra S. Haraguchi (also casual)
Kunio Hasebe
Joan Hironaka
Paul Jubinski
Lynette Y. I. Kam
James Karkheck
Mohammad M. Khan (also casual)
Stephen Y. Kuba
Alan H. Maruyama
Thomas M. Matsumoto (also casual)
Milton N. Matsuyama
Kenneth K. Obayashi
Kirit C. Patel
Jack A. Piper
Ronnie Roberts
Izumi Soma
Darlene I. Sugi
Doreen K. Suzuki
Neil S. Yamagata
Dave Y. Yonamine
Faith E. Yoshioka

Administrative Staff

Donald L. Dobbs
Jared S. Ikeda
Margaret M. Iwamoto
Doreen T. Leland
Jeanne H. Moriyasu
Evelyn G. Kwock Tasaka (also student help, casual)

Administrative Student Help

Lucille A. Barale
Alva A. Chai
Noreen P. Ettl
Sharen M. Gomoto
Faye M. Hinochi
Laurie L. Kino
Lily F. Ko
Susan Y. Koga
Carol Jane M. Maeda
Nancy Y. Nakagawa
Liana M. Petranek
Lani Rae S.L. Suiso
Beryl E. Tanaka
Carole A. Utsumi
JoAnn C. W. N. Wong